

Listing of Claims:

This listing of claims will replace all prior versions, and listings, of claims in the application:

1. (Canceled).
2. (Previously presented) A computer-implemented method for randomly walking through a hyper-text-linked document set comprising a plurality of documents, wherein at least a subset of the documents contain a plurality of links to other documents, each document being associated with a host, the method comprising:
 - a) selecting a host;
 - b) selecting at random a document associated with the host;
 - c) retrieving the selected document;
 - d) randomly choosing whether to select a random new document;
 - e) responsive to choosing to select the random new document:
 - e.1) selecting at random a new host from among the previously selected hosts;
 - e.2) selecting at random a new document associated with the new host; and
 - e.3) retrieving the selected new document;
 - f) responsive to choosing not to select the random new document:
 - f.1) selecting at random a link in the retrieved document; and
 - f.2) retrieving a document referenced by the selected link; and
 - g) repeating d), and then conditionally repeating e) or f) depending upon the choosing made in d).
3. (Canceled).
4. (Previously presented) The method of claim 2, wherein the document set is the World Wide Web, and wherein each document is a web page.

5. (Original) The method of claim 4, wherein each host corresponds to a domain.

6. (Previously presented) The method of claim 2, further comprising, concurrently with a) through d), with e) or f), and with g, performing a second two-level random walk through the hypertext-linked document set.

7. (Previously presented) A computer-implemented method for randomly walking through a hypertext-linked document set comprising a plurality of documents, wherein at least a subset of the documents contain a plurality of links to other documents, each document being associated with a host, the method comprising:

- a) initializing a host set;
- b) initializing a document set for each host in the host set;
- c) selecting at random a host from the host set;
- d) selecting at random a document from the document set of the selected host; and
- e) responsive to the selected document containing at least one link:
 - e.1) selecting at random a link from the selected document;
 - e.2) selecting a document corresponding to the selected link;
 - e.3) selecting a host corresponding to the selected document;
 - e.4) adding the selected host to the host set;
 - e.5) adding the selected document to the document set of the selected host; and
 - e.6) repeating e.1) through e.5) until all links have been traversed.

8. (Previously presented) The method of claim 7, wherein:
- e.4) is performed responsive to the selected host not being in the host set; and
 - e.5) is performed responsive to the selected document not being in the document set of the selected host.
9. (Previously presented) A computer-implemented method for randomly walking through a hypertext-linked document set comprising a plurality of documents, wherein at least a subset of the documents contain a plurality of links to other documents, each document being associated with a host, the method comprising:
- a) initializing a host set;
 - b) initializing a document set for each host in the host set;
 - c) selecting at random a host from the host set;
 - d) selecting at random a document from the document set of the selected host;
 - e) randomly choosing whether to select a random new document; and
 - f) responsive to choosing not to select a random new document and further responsive to the selected document containing at least one link:
 - f.1) selecting at random a link from the selected document;
 - f.2) selecting a document corresponding to the selected link;
 - f.3) selecting a host corresponding to the selected document;
 - f.4) adding the selected host to the host set;
 - f.5) adding the selected document to the document set of the selected host; and
 - f.6) repeating f.1) through f.5) until all links have been traversed.
10. (Canceled).

11. (Original) The method of claim 7, wherein the hypertext-linked document set is the World Wide Web, and wherein each document is a web page.

12. (Original) The method of claim 11, wherein each host corresponds to a domain.

13. (Previously presented) A computer-implemented method for measuring relative quality of a search engine index, comprising:

- a) performing a two-level random walk among documents within a document set, wherein at least a subset of the documents contain a plurality of links to other documents, each document being associated with a host, and wherein performing the two-level random walk comprises:
 - a.1) selecting a host;
 - a.2) selecting at random a document associated with the host;
 - a.3) retrieving the selected document;
 - a.4) selecting at random a link in the retrieved document;
 - a.5) retrieving a document referenced by the selected link; and
 - a.6) repeating a.4) and a.5) until all links have been traversed;
- b) for each document encountered in the random walk, determining whether the document is indexed by the search engine index; and
- c) aggregating the results of b).

14. (Canceled).

15. (Previously presented) A computer-implemented method for measuring relative quality of a search engine index, comprising:

- a) performing a two-level random walk among documents within a document set, by:
 - a.1) selecting a host;
 - a.2) selecting at random a document associated with the host;

- a.3) retrieving the selected document;
- a.4) randomly choosing whether to select a random new document;
 - a.4.1) responsive to choosing to select the random new document:
 - a.4.1.1) selecting at random a new host from among the previously selected hosts;
 - a.4.1.2) selecting at random a new document associated with the host; and
 - a.4.1.3) retrieving the selected new document;
 - a.4.2) responsive to choosing not to select the random new document:
 - a.4.2.1) selecting at random a link in the retrieved document; and
 - a.4.2.2) retrieving a document referenced by the selected link; and
- a.5) repeating a.4), and then conditionally repeating a.4.1) through a.4.1.3) or a.4.2) through a.4.2.2) depending upon the choosing made in a.4);
- b) for each document encountered in the random walk, determining whether the document is indexed by the search engine index; and
- c) aggregating the results of b).

16. (Previously presented) A computer-implemented method for measuring relative quality of a search engine index, comprising:

- a) performing a two-level random walk among documents within a document set, wherein at least a subset of the documents contain a plurality of links to other documents, each document being associated with a host, and wherein performing the two-level random walk comprises:
 - a.1) initializing a host set;

- a.2) initializing a document set for each host in the host set;
- a.3) selecting at random a host from the host set;
- a.4) selecting at random a document from the document set of the selected host;
- a.5) adding a host that is referenced by a selected link to the host set;
- a.6) adding a document referenced by the selected link to the document set of the selected host;
- a.7) responsive to the selected document containing at least one link:
 - a.7.1) selecting at random a link from the selected document;
 - a.7.2) selecting a document corresponding to the selected link;
 - a.7.3) selecting a host corresponding to the selected document; and
 - a.7.4) repeating a.5) through a.8) until all links have been traversed; and
- a.8) responsive to the selected document not containing at least one link, repeating a.3) through a.6), and further conditionally repeating a.7) or a.8), until all documents have been traversed;
- b) for each document encountered in the random walk, determining whether the document is indexed by the search engine index; and
- c) aggregating the results of b).

17. (Original) The method of claim 16, wherein:

- a.5) is performed responsive to the selected host not being in the host set; and
- a.6) is performed responsive to the selected document not being in the document set of the selected host.

18. (Original) The method of claim 13, wherein each document contains a plurality of words, and wherein b) comprises, for each document encountered in the random walk:

- b.1) selecting at least one word from the document;
- b.2) performing a query on the search engine index based on the selected at least one word, to obtain search results; and
- b.3) determining whether the document is included in the obtained search results.

19. (Original) The method of claim 18, wherein b.1) comprises selecting at least one word based on rarity.

20. (Canceled).

21. (Previously presented) A computer-implemented method for measuring relative quality of a target document in a document set comprising a plurality of documents, wherein at least a subset of the documents contain a plurality of links to other documents, the method comprising:

- a) performing a two-level random walk among documents within a document set; and
- b) determining a quality metric responsive to the number of documents encountered during the two-level random walk that link to the target document, the determining of the quality metric comprising determining a value for:

$$R(p) = d / T + (1 - d) \sum_{i=1}^k R(p_i) / C(p_i)$$

where:

$R(p)$ is the PageRank of target document p ;

$R(p_i)$ is the PageRank of document p_i ;

T is the total number of documents in the document set;

d is a damping factor such that $0 < d < 1$;

documents p_i, \dots, P_k each contain at least one link to target document p ; and
 $C(p_i)$ is the number of links out of document p_i .

22. (Previously presented) The method of claim 21, wherein b) comprises determining a quality metric responsive to the number of documents that link to the target document, and responsive to the quality metric of the linking documents.

23. (Canceled).

24. (Previously presented) A computer-implemented method for measuring relative quality of a target document in a document set comprising a plurality of documents, wherein at least a subset of the documents contain a plurality of links to other documents, wherein each document is associated with a host, the method comprising:

- a) performing a two-level random walk among documents within a document set by:
 - a.1) selecting a host;
 - a.2) selecting at random a document associated with the host;
 - a.3) retrieving the selected document;
 - a.4) randomly choosing whether to select a random new document;
 - a.5) responsive to choosing to select the random new document:
 - a.5.1) selecting at random a host from among the previously selected hosts;
 - a.5.2) selecting at random a document associated with the host; and
 - a.5.3) retrieving the selected document;
 - a.6) responsive to choosing not to select the random new document:

- a.6.1) selecting at random a link in the retrieved document;
and
 - a.6.2) retrieving a document referenced by the selected link;
and
 - a.7) repeating a.4), and then conditionally repeating a.5) through a.5.3) or a.6) through a.6.2) depending upon the choosing made in a.4); and
- b) determining a quality metric responsive to the number of documents encountered during the two-level random walk that link to the target document.

25. (Previously presented) A computer-implemented method for measuring relative quality of a target document in a document set comprising a plurality of documents, wherein at least a subset of the documents contain a plurality of links to other documents, wherein each document is associated with a host, the method comprising:

- a) performing a two-level random walk among documents within a document set, by:
 - a.1) initializing a host set;
 - a.2) initializing a document set for each host in the host set;
 - a.3) selecting at random a host from the host set;
 - a.4) randomly choosing whether to select a random new host;
 - a.5) responsive to choosing to select the random new host:
 - a.5.1) selecting at random a new host from among the previously selected hosts;
 - a.6) responsive to choosing not to select the random new host:
 - a.6.1) selecting at random a document from the document set of the selected host; and

- a.6.2) responsive to the selected document containing at least one link:
 - a.6.2.1) selecting at random a link from the selected document;
 - a.6.2.2) selecting a document corresponding to the selected link;
 - a.6.2.3) selecting a host corresponding to the selected document; and
 - a.6.2.4) adding the selected host to the host set;
 - a.6.2.5) adding the selected document to the document set of the selected host;
 - a.6.2.6) repeating a.6.2.1) through a.6.2.5) until all links have been traversed; and
 - a.7) repeating a.3) through a.4), and then conditionally repeating a.5) through a.5.1) or a.6) through a.6.2.6); and
 - b) determining a quality metric responsive to the number of documents encountered during the two-level random walk that link to the target document.
26. (Previously presented) The method of claim 21, further comprising:
- c) determining a quality metric for at least one additional target document; and
 - d) ranking the quality metric of the first target document with respect to the quality metrics of the additional target documents.

27. (Previously presented) A computer-implemented method for randomly walking through a hypertext-linked document set comprising a plurality of documents, wherein at least a subset of the documents contain a plurality of links to other documents, each document being associated with a host, the method comprising:

- a) selecting a host;

- b) selecting at random a document associated with the host;
- c) retrieving the selected document;
- d) randomly choosing whether to select a random new host;
- e) responsive to choosing to select the random new host:
 - e.1) selecting at random a new host from among the previously selected hosts; and
 - e.2) repeating b) through d), and then conditionally e) through e.2) or f) through f.3 until all documents have been traversed; and
- f) responsive to choosing not to select the random new host:
 - f.1) selecting at random a link in the retrieved document;
 - f.2) retrieving a document referenced by the selected link; and
 - f.3) repeating d), and then conditionally e) through e.2) or f) through f.3) until all links have been traversed.

28. (Previously presented) A computer-implemented method for measuring relative quality of a target document in a document set comprising a plurality of documents, wherein at least a subset of the documents contain a plurality of links to other documents, the method comprising:

- a) performing a two-level random walk among documents within a document set, by:
 - a.1) initializing a host set;
 - a.2) initializing a document set for each host in the host set;
 - a.3) selecting at random a host from the host set;
 - a.4) randomly choosing whether to select a random new host;
 - a.5) responsive to choosing to select a random new host:
 - a.5.1) selecting at random a new host from among the previously selected hosts;
 - a.6) responsive to choosing not to select the random new host:
 - a.6.1) selecting at random a document from the document set of the selected host; and

- a.6.2) responsive to the selected document containing at least one link:
 - a.6.2.1) selecting at random a link from the selected document;
 - a.6.2.2) selecting a document corresponding to the selected link;
 - a.6.2.3) selecting a host corresponding to the selected document; and
 - a.6.2.4) adding the selected host to the host set;
 - a.6.2.5) adding the selected document to the document set of the selected host;
 - a.6.2.6) repeating a.6.2.1) through a.6.2.5) until all links have been traversed; and
- a.7) repeating a.3) through a.4), and then conditionally repeating a.5) through a.5.1) or a.6) through a.6.2.6); and
- b) determining a quality metric responsive to the number of documents encountered during the two-level random walk that link to the target document;
- c) determining a quality metric for at least one additional target document; and
- d) ranking the quality metric of the first document with respect to the quality metrics of the additional target documents.

29. (Canceled).

30. (Previously presented) A computer program product comprising a computer-readable medium having computer-readable code embodied therein for randomly walking through a hypertext-linked document set comprising a plurality

of documents, wherein at least a sub-set of the documents contain a plurality of links to other documents, each document being associated with a host, the computer program product comprising:

- a) computer-readable program code devices configured to cause a computer to select a host;
- b) computer-readable program code devices configured to cause a computer to select at random a document associated with the host;
- c) computer-readable program code devices configured to cause a computer to retrieve the selected document;
- d) computer-readable program code devices configured to cause a computer to randomly choose whether to select a random new document;
- e) computer-readable program code devices configured to cause a computer to, responsive to choosing to select the random new document:
 - e.1) select at random a new host from among the previously selected hosts; and
 - e.2) select at random a new document associated with the host; and
 - e.3) retrieve the selected new document;
- f) computer-readable program code devices configured to cause a computer to, responsive to choosing not to select the random new document:
 - f.1) select at random a link in the retrieved document; and
 - f.2) retrieve a document referenced by the selected link; and
- g) computer-readable program code devices configured to cause a computer to repeat the operations of d) and then conditionally repeat the operations of e) or f) depending on the choice made in d).

31. (Canceled).

32. (Previously presented) The computer program product of claim 30, wherein the document set is the World Wide Web, and wherein each document is a web page.

33.-34 (Canceled).

35. (Previously presented) A computer program product comprising a computer-readable medium having computer-readable code embodied therein for randomly walking through a hypertext-linked document set comprising a plurality of documents, wherein at least a subset of the documents contain a plurality of links to other documents, each document being associated with a host, the computer program product comprising:

- a) computer-readable program code devices configured to cause a computer to initialize a host set;
- b) computer-readable program code devices configured to cause a computer to initialize a document set for each host in the host set;
- c) computer-readable program code devices configured to cause a computer to select at random a host from the host set;
- d) computer-readable program code devices configured to cause a computer to select at random a document from the document set of the selected host; and
- e) computer-readable program code devices configured to cause a computer to, responsive to the selected document containing at least one link:
 - e.1) select at random a link from the selected document;
 - e.2) select a document corresponding to the selected link;
 - e.3) select a host corresponding to the selected document; and
 - e.4) add the selected host to the host set;
 - e.5) add the selected document to the document set of the selected host; and

e.6) repeat the operations of e.1) through e.5) until all links have been traversed.

36. (Original) The computer program product of claim 35, wherein:
the computer-readable program code devices configured to cause a computer to add the selected host to the host set operate responsive to the selected host not being in the host set; and
the computer-readable program code devices configured to cause a computer to add the selected document to the document set of the selected host operate responsive to the selected document not being in the document set of the selected host.

37. (Previously presented) A computer program product comprising a computer-readable medium having computer-readable code embodied therein for randomly walking through a hypertext-linked document set comprising a plurality of documents, wherein at least a sub-set of the documents contain a plurality of links to other documents, each document being associated with a host, the computer program product comprising:

- a) computer-readable program code devices configured to cause a computer to initialize a host set;
- b) computer-readable program code devices configured to cause a computer to initialize a document set for each host in the host set;
- c) computer-readable program code devices configured to cause a computer to select at random a host from the host set;
- d) computer-readable program code devices configured to cause a computer to select at random a document from the document set of the selected host;
- e) computer-readable program code devices configured to cause a computer to randomly choose whether to select a random new document; and

- f) computer-readable program code devices configured to cause a computer to, responsive to choosing not to select a random new document, and further responsive to the selected document containing at least one link:
 - f.1) select at random a link from the selected document;
 - f.2) select a document corresponding to the selected link;
 - f.3) select a host corresponding to the selected document; and
 - f.4) add the selected host to the host set;
 - f.5) add the selected document to the document set of the selected host; and
 - f.6) repeat the operations of f.1 through f.5) until all links have been traversed.

38. (Canceled).

39. (Original) The computer program product of claim 35, wherein the hypertext-linked document set is the World Wide Web, and wherein each document is a web page.

40. (Original) The computer program product of claim 39, wherein each host corresponds to a domain.

41. (Previously presented) A computer program product comprising a computer-readable medium having computer-readable code embodied therein for measuring relative quality of a search engine index, the computer program product comprising:

- a) computer-readable program code devices configured to cause a computer to perform a two-level random walk among documents within a document set, wherein at least a subset of the documents contain a plurality of links to other documents, each document being associated with a host, and wherein the computer-readable

program code devices configured to cause a computer to perform a two-level random walk comprise:

- a.1) computer-readable program code devices configured to cause a computer to select a host;
- a.2) computer-readable program code devices configured to cause a computer to select at random a document associated with the host;
- a.3) computer-readable program code devices configured to cause a computer to retrieve the selected document;
- a.4) computer-readable program code devices configured to cause a computer to select at random a link in the retrieved document;
- a.5) computer-readable program code devices configured to cause a computer to retrieve a document referenced by the selected link; and
- a.6) computer-readable program code devices configured to cause a computer to repeat the operations of a.4) and a.5) until all links have been traversed;
- b) computer-readable program code devices configured to cause a computer to, for each document encountered in the random walk, determine whether the document is indexed by the search engine index; and
- c) computer-readable program code devices configured to cause a computer to aggregate the results of the operations of b).

42.-43. (Canceled).

44. (Previously presented) A computer program product comprising a computer-readable medium having computer-readable code embodied therein for measuring relative quality of a search engine index, the computer program product comprising:

- a) computer-readable program code devices configured to cause a computer to perform a two-level random walk among documents within a document set, wherein at least a subset of the documents contain a plurality of links to other documents, each document being associated with a host, and wherein the computer-readable program code devices configured to cause a computer to perform a two-level random walk comprise:
 - a.1) computer-readable program code devices configured to cause a computer to initialize a host set;
 - a.2) computer-readable program code devices configured to cause a computer to initialize a document set for each host in the host set;
 - a.3) computer-readable program code devices configured to cause a computer to select at random a host from the host set;
 - a.4) computer-readable program code devices configured to cause a computer to select at random a link from a document in the document set of the selected host;
 - a.5) computer-readable program code devices configured to cause a computer to add a host referenced by the link to the host set;
 - a.6) computer-readable program code devices configured to cause a computer to add a document referenced by the link to the document set of the selected host;
 - a.7) computer-readable program code devices configured to cause a computer to, responsive to the selected document containing at least one link:

- a.7.1) select at random a link from the selected document;
- a.7.2) select a document corresponding to the selected link;
- a.7.3) select a host corresponding to the selected document;
- and
- a.7.4) repeat the operations of a.5) through a.8) until all links have been traversed; and
- a.8) computer-readable program code devices configured to cause a computer to, responsive to the selected document not containing at least one link, repeat the operations of a.3) through a.6), and further conditionally repeating a.7) or a.8), until all documents have been traversed;
- b) computer-readable program code devices configured to cause a computer to, for each document encountered in the random walk, determine whether the document is indexed by the search engine index; and
- c) computer-readable program code devices configured to cause a computer to aggregate the results of the operations of b).

45. (Original) The computer program product of claim 44, wherein:
- the computer-readable program code devices configured to cause a computer to add the selected host to the host set are configured to cause a computer to add the selected host responsive to the selected host not being in the host set; and
 - the computer-readable program code devices configured to cause a computer to add the selected document to the document set of the selected host are configured to cause a computer to add the selected document responsive to the selected document not being in the document set of the selected host.

46. (Previously presented) The computer program product of claim 41, wherein each document contains a plurality of words, and wherein the computer-

readable program code devices configured to cause a computer to determine whether the document is indexed by the search engine index comprise computer-readable program code devices configured to, for each document encountered in the random walk:

- b.1) select at least one word from the document;
- b.2) perform a query on the search engine index based on the selected at least one word, to obtain search results; and
- b.3) determine whether the document is included in the obtained search results.

47. (Original) The computer program product of claim 46, wherein the computer-readable program code devices configured to select at least one word from the document comprise computer-readable program code devices configured to select at least one word based on rarity.

48. (Canceled).

49. (Previously presented) A computer program product comprising a computer-readable medium having computer-readable code embodied therein for measuring relative quality of a target document in a document set comprising a plurality of documents, wherein at least a subset of the documents contain a plurality of links to other documents, the computer program product comprising:

computer-readable program code devices configured to cause a computer to perform a two-level random walk among documents within a document set; and

computer-readable program code devices configured to cause a computer to determine a quality metric responsive to the number of documents encountered during the two-level random walk that link to the target document and to further determine a value for:

$$R(p) = d / T + (1 - d) \sum_{i=1}^k R(p_i) / C(p_i)$$

where:

$R(p)$ is the PageRank of target document p ;

$R(p_i)$ is the PageRank of document p_i ;

T is the total number of documents in the document set;

d is a damping factor such that $0 < d < 1$;

documents p_1, \dots, p_k each contain at least one link to target document p ; and

$C(p_i)$ is the number of links out of document p_i .

50. (Previously presented) The computer program product of claim 49, wherein the computer-readable program code devices configured to cause a computer to determine a quality metric comprise computer-readable program code devices configured to cause a computer to determine a quality metric responsive to the number of documents that link to the target document, and responsive to the quality metric of the linking documents.

51. (Canceled).

52. (Previously presented) A computer program product comprising a computer-readable medium having computer-readable code embodied therein for measuring relative quality of a target document in a document set comprising a plurality of documents, wherein at least a subset of the documents contain a plurality of links to other documents, and wherein each document is associated with a host, the computer program product comprising:

computer-readable program code devices configured to cause a computer to perform a two-level random walk among documents within a document set, by:

- a.1) selecting a host;
- a.2) selecting at random a document associated with the host;
- a.3) retrieving the selected document;

- a.4) randomly choosing whether to select a random new document;
- a.5) responsive to choosing to select the random new document:
 - a.5.1) selecting at random a host from among the previously selected hosts; and
 - a.5.2) selecting at random a document associated with the host; and
 - a.5.3) retrieving the selected document;
- a.6) responsive to choosing not to select the random new document:
 - a.6.1) selecting at random a link in the retrieved document; and
 - a.6.2) retrieving a document referenced by the selected link; and
- a.7) repeating the operations of a.4), and then conditionally repeating the operations of a.5) through a.5.3) or a.6) through a.6.2) depending upon the choosing made in a.4); and

computer-readable program code devices configured to cause a computer to determine a quality metric responsive to the number of documents encountered during the two-level random walk that link to the target document.

53. (Previously presented) A computer program product comprising a computer-readable medium having computer-readable code embodied therein for measuring relative quality of a target document in a document set comprising a

plurality of documents, wherein at least a subset of the documents contain a plurality of links to other documents, wherein each document is associated with a host, the computer program product comprising:

computer-readable program code devices configured to cause a computer to perform a two-level random walk among documents within a document set, by:

- a.1) initializing a host set;
- a.2) initializing a document set for each host in the host set;
- a.3) selecting at random a host from the host set;
- a.4) randomly choosing whether to select a random new host;
- a.5) responsive to choosing to select the random new host:
 - a.5.1) selecting at random a host from among the previously selected hosts;
- a.6) responsive to choosing not to select the random new host:
 - a.6.1) selecting at random a document from the document set of the selected host;
 - a.6.2) adding the selected host to the host set;
 - a.6.3) adding the selected document to the document set of the selected host;
 - a.6.4) responsive to the selected document containing at least one link:
 - a.6.4.1) selecting at random a link from the selected document;
 - a.6.4.2) selecting a document corresponding to the selected link;
 - a.6.4.3) selecting a host corresponding to the selected document; and
 - a.6.4.4) repeating the operations of a.6.2) through a.6.4.3) until all links have been traversed; and

- a.7) repeating the operations of a.3) through a.4), and then conditionally repeating a.5) through a.5.1) or a.6) through a.6.4.4; and

computer-readable program code devices configured to cause a computer to determine a quality metric responsive to the number of documents encountered during the two-level random walk that link to the target document.

54. (Previously presented) The computer program product of claim 49, further comprising:

- c) computer-readable program code devices configured to cause a computer to determine a quality metric for at least one additional target document; and
- d) computer-readable program code devices configured to cause a computer to rank the quality metric of the first target document with respect to the quality metrics of the additional target documents.

55. (Previously presented) A computer program product comprising a computer-readable medium having computer-readable code embodied therein for randomly walking through a hypertext-linked document set comprising a plurality of documents, wherein at least a subset of the documents contain a plurality of links to other documents, each document being associated with a host, the computer program product comprising:

- a) computer-readable program code devices configured to cause a computer to select a host;
- b) computer-readable program code devices configured to cause a computer to select at random a document associated with the host;
- c) computer-readable program code devices configured to cause a computer to retrieve the selected document;

- d) computer-readable program code devices configured to cause a computer to randomly choose whether to select a random new host;
- e) computer-readable program code devices configured to cause a computer to, responsive to choosing to select the random new host:
 - e.1) select at random a new host from among the previously selected hosts; and
 - e.2) repeat the operations of b) through d) and then conditionally e) through e.2) or f) through f.3) until all documents have been traversed; and
- f) computer-readable program code devices configured to cause a computer to, responsive to choosing not to select the random new host:
 - f.1) select at random a link in the retrieved document;
 - f.2) retrieve a document referenced by the selected link; and
 - f.3) repeat the operations of d), and then conditionally e) through e.2) or f) through f.3) until all links have been traversed.

56. (Previously presented) A computer program product comprising a computer-readable medium having computer-readable code embodied therein for measuring relative quality of a target document in a document set comprising a plurality of documents, wherein at least a subset of the documents contain a plurality of links to other documents, the computer program product comprising:

- a) computer-readable program code devices configured to cause a computer to perform a two-level random walk among documents within a document set by:
 - a.1) initializing a host set;
 - a.2) initializing a document set for each host in the host set;
 - a.3) selecting at random a host from the host set;
 - a.4) randomly choosing whether to select a random new host;

- a.5) responsive to choosing to select a random new host:
 - a.5.1) selecting at random a new host from among the previously selected hosts;
- a.6) responsive to choosing not to select the random new host:
 - a.6.1) selecting at random a link from a document in the document set of the selected host;
 - a.6.2) adding the host referenced by the link to the host set;
 - a.6.3) adding the document referenced by the link to the document set of the selected host;
 - a.6.4) responsive to the selected document containing at least one link:
 - a.6.4.1) selecting at random a link from the selected document;
 - a.6.4.2) selecting a document corresponding to the selected link;
 - a.6.4.3) selecting a host corresponding to the selected document;
 - a.6.4.4) repeating the operations of a.6.2) through a.6.4.3) until all links have been traversed; and
- a.7) responsive to the selected document not containing at least one link, repeating the operations of a.3) through a.4), and then conditionally repeating a.5) through a.5.1) or a.6) through a.6.4.4);
- b) computer-readable program code devices configured to cause a computer to determine a quality metric responsive to the number of documents encountered during the two-level random walk that link to the target document;
- c) computer-readable program code devices configured to cause a computer to determine a quality metric for at least one additional target document; and

- d) computer-readable program code devices configured to cause a computer to rank the quality metric of the first document with respect to the quality metrics of the additional target documents.

57. (Previously presented) A system for randomly walking through a hypertext-linked document set comprising a plurality of documents, wherein at least a subset of the documents contain a plurality of links to other documents, each document being associated with a host, the system comprising:

- a) a host selector;
- b) a random document selector, coupled to the host selector, for selecting at random a document associated with the host;
- c) a document retriever, coupled to the random document selector, for retrieving the selected document; and
- d) a link selector, coupled to the document retriever;

wherein, responsive to the host selector randomly choosing to select a random host:

the host selector selects at random a host from among the previously selected hosts;

the random document selector selects at random a document associated with the host; and

the document retriever retrieves the selected document; and

wherein, responsive to the host selector randomly choosing not to select a random host:

the link selector selects at random a link in the retrieved document; and

the document retriever retrieves a document referenced by the selected link; and

wherein the link selector, the random document selector, and the document retriever repeat their respective operations until all links have been traversed.

58.-62. (Canceled).

63. (Previously presented) The method of claim 2, wherein the repeating of d) and the conditional repeating of e) or f) continues until all documents have been traversed.

64. (Previously presented) The method of claim 7, further comprising:
f) repeating c) through d), and further conditionally repeating e) if the selected document contains at least one link, until all documents have been traversed.

65. (Previously presented) The method of claim 9, further comprising:
g) repeating c) through e), and further conditionally repeating f) if a random new document is not chosen, until all documents have been traversed.

66. (Previously presented) The method of claim 15, wherein the repeating of a.4) and the conditional repeating of a.4.1) through a.4.1.3) or a.4.2) through a.4.2.2) continues until all documents have been traversed.

67. (Previously presented) The method of claim 24, wherein the repeating of a.4) and the conditional repeating of a.5) through a.5.3) or a.6) through a.6.2) continues until all documents have been traversed.

68. (Previously presented) The method of claim 25, wherein the repeating of a.3) through a.4) and the conditional repeating of a.5) through a.5.1) or a.6) through a.6.2.6) continues until all documents have been traversed.

69. (Previously presented) The method of claim 28, wherein the repeating of a.3) through a.4) and the conditional repeating of a.5) through a.5.1 or a.6) through a.6.2.6) continues until all documents have been traversed.

70. (Previously presented) The computer program product of claim 30, wherein the computer-readable program code devices are further configured to continue to cause a computer to repeat a.4) and the conditionally repeat a.4.1) through a.4.1.3) or a.4.2) through a.4.2.2) until all documents have been traversed.

71. (Previously presented) The computer program product of claim 35, further comprising:

- f) computer readable program code devices configured to cause a computer to repeat c) through d), and further conditionally repeat e) if the selected document contains at least one link, until all documents have been traversed.

72. (Previously presented) The computer program product of claim 37, further comprising:

- g) computer readable program code devices configured to cause a computer to repeat c) through e), and further conditionally repeat f) if a random new document is not chosen, until all documents have been traversed.

73. (Previously presented) The computer program product of claim 52, wherein the repeating of a.4) and the conditional repeating of a.5) through a.5.3) or a.6) through a.6.2) continues until all documents have been traversed.

74. (Previously presented) The computer program product of claim 53, wherein the repeating of a.3) through a.4) and the conditional repeating of a.5) through a.5.1) or a.6) through a.6.4.4) continues until all documents have been traversed.

75. (Previously presented) The computer program product of claim 56, wherein the repeating of a.3) through a.4) and the conditional repeating a.5) through a.5.1) or a.6) through a.6.4.4) continues until all documents have been traversed.